Navy Regional Data Automation Center

Norfolk, Virginia

"EVERYTHING YOU EVER WANTED TO KNOW ABOUT MICROCOMPUTERS"

(But didn't know WHO to ask)

An informal, but detailed look into
current microcomputer technology, trends,
standards, pitfalls and recommendations.

by

ETC Leland W. Slater
Microsystems Evaluator
Micro/Mini Evaluation Team

1 June 1982

So, you're off to select and buy your command's computer. You've been thoroughly amazed at what microcomputers can do. You've seen them calculate formulas, balance checkbooks, predict trends, perhaps you've even seen a data management system in use. With all of those capabilities available on most micros, your can't really make a wrong selection - or can you?

The selection of a microcomputer can be very intimidating. One earful of a salesperson's golden words, or a friend extolling the virtues of his personal micro, and you are left with the impression that all other brands are left in the dust. Problem is, the next guy you talk to recommends something else. And why can't you take that neat program on the first machine and use it on the second? Isn't there someone who will give you a straight unbiased opinion? Step into my office and let's talk about it.

Because this is an article and not a book, let me eliminate a few "sticky" areas first. A small percentage of users want an inexpensive system to use for games, as a "desktop secretary", for plotting stock trends, Computer Assisted Instruction (CAI), and other "consumer" type activities. Todays Personal Computers (PCs) excel in these types of areas, and I would refer you to your local computer store to look at the Apple, Atari, IBM, Radio Shack, Commodore, and similar machines.

Another small percentage of users need the capability of having multiple people access one or more databases simultaneously. If you are in this catagory, I suggest you get in touch with the NARDAC Norfolk Micro Team directly, because you need to look into a LOT of different areas, and you may need lots of handholding.

If you fall into neither of the above catagories (as the vast majority of serious users will) then grab a cup of coffee and a pad of paper, and let's figure out what you need.

We will eventually make recommendations at two levels. The first level we will call a "requirement". This is the absolute minimum you should include in your specification. To NOT include this item could have potentially grave repercussions, usually resulting from total incompatibility or serious inadequacy. As a minimum, you will experience tremendous inconvenience. We would NOT recommend that you forgo ANY item listed as a "requirement" without consulting directly with the Micro Evaluation Team. The second level will be called "recommendations", and will include things you should ask for if at all possible, or if conditions for that requirement exist as explained in the description which preceeds each reqirement and/or recommendation. Some items will have only a requirement, some will have only a recommendation. Most items have both.

When you select a computer, you MUST remain impartial. If you have a machine in mind when you start out, you will not be selecting a machine by its merits, but rather justifying why you can live with its various weaknesses. ALL microcomputers have their strong points and weak points. The object is to filter through YOUR requirements, separating the wheat from the chaff, until a clear set of minimum requirements is defined, as well as a list of "nice to have" optional features. What do you really need, and what can you live without.

Let's pause a moment so you can get a feel for your surroundings. The computer category you are in can be called the "serious, medium scale small business" group. You need to get fast results from medium to large databases, and manipulate words quickly and efficiently. Your objective is to do more work in less time and with less effort. You may not be profit oriented, but you are very definitely running a business. Supply department becomes a parts store, mess halls become restaurants. Hey, that's something important. We need lots of really SERIOUS software. Not mom and pop size. You're operations are fast paced and time is critical. Your work hours, perhaps even your life, may depend on selecting a proper computer. "Give me liberty or give me death". Snappy phrase. May even become famous someday.

Perhaps the most important difference between you and a commercial firm is that you are part of a huge team. In a business, it is not normally profitable or wise to give your competitors things like programs/applications you have developed and data bases you have collected. The military is just the opposite. We can derive great value from sharing. There is also a great personal satisfaction when you share your efforts with your counterparts at other commands. It may even be a great "cumshaw" trading medium. "You give me your Pre-Expended Bin application and I'll give you our COSAL database." Ah-Ha! Sounds like something important. We had better make sure we can talk to anybody else out there, even though they may all have different machines.

Talking about different machines, is your application so limited and specialized that one, and only one manufacturer's machine is capable of performing it? Probably not. If it is, then you have no need to read further. Your time would be better spent thinking of ways to proceed with the sole-source procurement paper mill, and in justifying to yourself and others why you are pushing a brand name instead of a requirement. But lets assume you are still requirements oriented. You are going to want to select your system from the broadest, and therefore the most competitive, product lines available. You want freedom to generically define both the hardware AND software to be used on your system.

Well, here we are. We haven't even BEGUN to talk about YOUR actual application, have we? There's a method to this madness, however. In the violently changing world of the microcomputer, the only thing you can be sure of is change. There IS, hovewer, an oasis of stability. There IS a way to have access to all the high performance software you're ever likely to need. You CAN select your system and be assured that you can share applications with the maximum number of other computer brands. And most important, you can do it by PURE SPECIFICATION. That should make the pork chops happy.

THE OPERATING SYSTEM

The heart of any computer, from the user's standpoint, is the "Operating System", or OS. Remembering that we are talking about SERIOUS applications and speed, we shall modify the phrase to Disk Operating System", or DOS. The DOS controls virtually all aspects of a computer. Indeed, it gives a computer its "personality". A DOS is closely associated with the storage medium used by a computer. It defines how data is written onto a disk, as well as the physical charactistics of the medium. The DOS also acts as a "buffer", separating the application being run from "grudge work" unique to the hardware it's running on. An application program may say to the DOS, "Save the stuff in memory from 'here' to 'there' to disk file called 'BLOB'". The DOS will then do all the dirty work necessary to complete the task.

A DOS, for our purposes, should have the maximum amount of SERIOUS business software available, so that we may choose from the broadest range of products. The DOS should run on a variety of hardware - NOT just that of a particular manufacturer. It should also provide a common exchange format, so that the application or database you develop can be transferred between manufacturers conveniently and reliably. Naturally, the DOS should be versatile, well tested, bug free, inexpensive and widely accepted by the industry. What we have just described is a family of operating systems which are "CP/M Compatible".

The CP/M operating system (Control Program for Microcomputers) has been around since 1975. It was developed by Gary Killdall, a Naval officer at the Naval Post Graduate School in Monterey, for the Intel MDS development system. To make a long story short, Intel didn't want it. Gary beefed it up a little, started Digital Research Inc, and began practically GIVING away version 1.3. It was just what the industry had been waiting for. The timing could not have been better. Version 1.4 sold hundreds of thousands of copies at the ridiculously low price of $150.00 - low enough for most manufacturers to give it away with their system for free. Hundreds of companies were writing CP/M compatible software, assured that their product would run on machines from hundreds of different manufacturers. Other companies began writing CP/M compatible operating systems as well. Infosoft (IOS) and Cromemco (CDOS) gave healthy competition to Digital Research almost immediately. Today, several other manufacturers make compatible operating systems.

All CP/M compatible programs and operating systems can exchange programs and data via the 8 inch Single Sided, Single Density (SSSD) soft sectored disk format - the ONLY such non-proprietary capability in the industry. Today, CP/M 2.2 compatible products can lay claim to more serious business software than any other proprietary or non-proprietary micro operating system available. There are, conservatively, over a half million users of CP/M compatible products. CP/M compatibility allowed hardware manufacturers to increase their system performance without requiring new software, and it allowed software vendors to increase software performance without worrying about the hardware it would run on. CP/M runs on everything from the $99.00 Sinclair ZX-81 to the Digital Equipment Corp (DEC) VAX minicomputer under emulation. THAT'S flexibility! It also runs on hardware from companies who traditionally push their own proprietary operating systems, companies such as IBM, XEROX, Radio Shack, Apple, Commodore, Heath/Zenith, Lexitron, Datapoint, Burroughs, Monroe, Wang, CPT, DEC, and others.

There must be SOMETHING awfully good going for CP/M, and there is. Remember, when a manufacturer blasts petty problems of CP/M and claims that his proprietary DOS is better, remember - better for WHOM?. Better for HIM, that's who. He would just love to lock you into his product line until you cannot afford to change brands to upgrade or expand. Software and transportability is the key to buying computers today. Hardware is almost insignificant by comparison. So, here is our first "standard":

REQUIREMENT: System must run a CP/M 2.2 compatible operating system.

Okay. There it is. Our big Micro Team prejudice, and all the reasons behind it. We now have a foundation on which to build your system. What are some of the other "standards" that go along with CP/M?

MICROPROCESSOR TYPE

CP/M (read "CP/M compatible operating systems") can run on any machine which has, or can emulate, the instruction set of the 8080 microprocessor. These include the popular Z-80 and 8085 processors, both of which are supersets of the 8080. The Z-80 has twice as many registers as the 8080, twice as many instructions, and normally runs twice as fast. The vast majority of today's hardware products use the Z-80. Software written for the Z-80 takes less space and has better performance. ALL 8080 software can run on the Z-80, but Z-80 software cannot run on the 8080. Not using a Z-80 compatible CPU will reduce software selection somewhat, and will reduce the performance of any package.

A special caveat should be noted here regarding some of the new 16 bit microprocessors which can run "CP/M", most notably the 8086 CPU running CPM-86. The vast majority of CP/M software - especially the high performance packages such as word processors, applications generators and electronic spreadsheets - are written in 8080 or Z-80 machine code. These packages can NOT be run under CP/M-86. What this means to you, in a nutshell, is that nearly all of the software that runs under normal CP/M compatible operating systems is NOT available under CPM-86! This void will probably be filled adequately, in time, but that doesn't do you any good right now. A few emulators are available which simulate the 8080 on the 8086, Z-8000 or 68000, but they are SLOW, SLOW, SLOW! If you feel you MUST get a 16 bit processor (you probably don't need it), specify a "dual" processor - there are several available. These products usually contain one 8 bit processor and one 16 bit processor, and allow you to run programs using either CPU under different operating systems. Some implementations allow you to operate two operating systems at once. At least one even allows you to mix programs on a common disk, and the machine figures out which processor to use by looking at the code to be executed! Dual processors are available using combinations of 8080/8086, 8085/8088, Z-80/8086, Z-80/Z-8000, Z-80/68000, Z-80/Pascal Microengine and others. Most are available only to those using a "system buss", which will be discussed later.

Remember, however, you probably DON'T really need a 16 bit processor for your requirements. A dual processor system, while better than a 16 bit alone, will add cost and complexity to your system, and reduce competitive bid prospects. If you are CERTAIN that you need a 16 bit CPU (as is the case when you have a piece of unique software which only runs under that CPU's opcodes and operating system, or if you're into HEAVY number crunching) then go ahead and specify a dual processor. If you are not sure what you'll need, specify the Z-80 and a buss design so that you may upgrade to a dual processor later, should it become necessary.

REQUIREMENT: Processor must execute the 8080 CPU instruction set.

RECOMMENDATION: Z-80 compatible CPU. 4MHZ speed preferred, 2 or 6 MHZ speeds optional. Dual processors are recommended if the need for a 16 bit CPU is absolutely required.

SYSTEM MEMORY

Every computer uses internal memory to execute programs. This memory is called "Random Access Memory", or RAM. RAM is ridiculously cheap nowadays, and there is no reason why every CP/M machine can't have a full 64K of RAM. So we just specify 64K of RAM and our troubles are over, right? Well, not exactly. The 64K address space contains all the software and data areas required to run your system. As a minimum, this includes the 6 to 8K for CP/M itself. The typical CP/M system has at least 56K actually free for the user to execute his programs in. This area is referred to as the "Transient Program Area", or TPA. There are many software products, especially high performance packages, which require the full 56K minimum TPA to operate. Most packages do not require the 56K, but rather use a second threshold of 48K. Nearly ALL serious software requires a minimum 48K TPA. This all seems pretty clear cut, except that manufactures love to tell of their TOTAL RAM capacity, instead of their TPA size. A system with "64K of RAM", for instance, may actually have 8K for CP/M and 8K for a special memory mapped video display. The user is restricted to 48K programs.

There are many other advantages to having a maximum amount of TPA. Nearly all commercial software packages use whatever memory is left over after their program is loaded to hold work buffers, indexes, etc. There is often a VAST improvement in performance between a package executing on a 48K TPA machine and the some package running in a 56K TPA, due to the larger area used for searches, sorts and indexes. Word processors experience a dramatic performance boost as well, since more pages can be in memory at once. In addition, there are fewer disk accesses, since the larger work areas mean more data is stored/retrieved with each access. This reduces the number of times the disk has to search for the location of disk information, as well a the time required to read or write a given amount of data. Word processors, application generators and electronic spread sheets - the very programs you will most likely be using - are the ones which benefit the most from a large TPA. Some manufacturers have even devised sneaky ways of increasing the TPA up to 62K or more by placing most of CP/M up above the normal limit of 64K, sometimes along with highspeed "track buffers" (Cache memory) to speed disk access. These systems provide the highest possible performance available

today under CP/M, often having an order of magnitude in speed increase over more conventional arrangements. Specification of TPA's larger than 56K, or for Cache memory, will drastically reduce vendor possibilities for competitive bid.

REQUIREMENT: Minimum 48K TPA.

RECOMMENDED: Minimum 56K TPA (You may wish to specify Cache memory for time critical, disk intensive applications).

VIDEO TERMINAL

Before we move on to things which are "nice to have" for various reasons, let's touch on one more item considered standard in the serious business software community. The video display terminal, referred to as the "CRT", is the primary method for operator control and interaction. Virtually all CP/M software expects an 80 character by 24 line CRT as a minimum. There are many reasons for this, not the least of which is the fact that mini and maxi computers have used 80 x 24 CRT's since before the micro was just a twinkle in an Intel engineer's eye. There are very few terminals on the market today which are not 80 x 24, and nearly all of those are 132 columns.

The 80 x 24 format was a compromise between the largest amount of data that could be placed on a display, and the most readable character size. On a 12" display, the density and character size of the 80 x 24 format is just about ideal. The characters are dense enough to allow an adequate amount of data within the users field of vision, and can display the full width of a standard document (70 to 78 chars). Increase the density and eyestrain can become a factor. A green (or amber) character display also helps reduce eyestrain. Character densities above 85 chars per line and 26 lines high are normally NOT an advantage unless a requirement exists for super dense video displays which cannot be met with hardcopy printouts.

For obvious reasons, the keyboard should be a standard "QWERTY" style and have a shift key (don't laugh!). In addition, the "control" key is universally used for special command entries. Beyond this, things such as alpha/caps lock, cursor controls, numeric keypad and function keys may be specified.

In any event, our concern is to be compatible with CP/M software, and it (along with everyone else) expects at least an 80 x 24 CRT. So...

REQUIREMENT: DISPLAY: 80 character (minimum) by 24 line (minimum) CRT, upper & lower case chars.
KEYBOARD: "QWERTY" style, shift key, control key.

RECOMMENDED: DISPLAY: A 12" physical display size is optimum (others may be used if physical constraints dictate). Green or amber chars on black background are preferred.
KEYBOARD: Should have Shift key, alpha lock and "control" keys.

DISK DRIVES

Up until now we have given requirements which will allow virtually all CP/M manufacturers to bid, and recommendations which, if all are specified, will allow perhaps 85% to participate. Now, sadly, we come to an item that some manufacturers have used to isolate themselves from everyone else.

Disk drives are used to hold programs, databases and documents. They are used as working storage during sorts and merges, and for data archiving ("off-line" storage). But perhaps the most important use, for the Navy in particular, is for transporting data and programs between systems.

Any disk drive used should be capable of holding the required amount of data. This can be easily found be multiplying the number of characters required for each record by the maximum number of records in the database. Calculate this figure for each datafile which must be in use simultaneously. Include any expected growth factors. Add the databases together and add a 20 to 50% fudge factor - depending on how confident you are of your estimates.

For example, let's take a maintenance application. The first database you may need is one which describes equipment - noun name, serial number, APL, FSN, SNAFU status, person responsible, and remarks. Let's say all the fields added together come to 80 characters (just like a punched card - gag, retch). And lets say we have 1500 pieces of equipment total. Thats 1500 x 80, or 120K. Our second data base contains our CSMP, and contains 150 - 250 compressed 2 - kilos (maintenance documents) of 3000 characters each. Thats 3000 x 250, or 750K. We give a fudge factor of 35% because we know the numbers are fairly stable, but could grow under casuality conditions. This gives us (120 + 750) * 1.35 = 1174.5K total - 1.2Mb for argument's sake. Funny, thats about how much storage is available on the typical Double Sided, Double Density (DSDD) 8 inch disk drive. Must have been a coincidence. Storage capacities for other 8 inch formats are approximately 500K for Single Sided, Double Density (SSDD), and 241K for Single Sided, Single Density (SSSD). You may recall we mentioned the 8 inch SSSD format as being a universal CP/M transport medium. All 8 inch CP/M drives, regardless of their maximum density format, can also read, write and use the standard SSSD CP/M interchange format (241K/disk). This is not the case for 5 inch (actually 5 1/4 inch) drives. Remember that as we talk about what seems to be a recent sad trend in the industry.

Five inch drives have been around a long time. They are a favorite in dedicated word processors where space is at a premium, cost is the driving concern, and no one ever thinks about exchanging diskettes. As their number grew, so did all the unique ways of formatting them. Of the dozens of 5 inch formats used under CP/M today, there are only a few that can talk with any one but themselves. Any passing of data between different manufactures has to be done via phone line or physical (hardwired) data link. We receive frequent calls asking how to send data from one 5 inch disk to another. We usually recommend a modem hookup and offer one of our public domain communications programs - if they can figure a way to get it off our standard 8 inch disk!

Naturally, the manufacturers will attest to the fact that not only is their new 5 inch format better than the standard 8 inch, but "better than any other 5 inch because _____." (Fill in the blank). Again comes the question - better for WHOM?

This time, there may be some bonafide cases to forego transportability - especially if physical size is to be minimized. Cost is usually NOT a valid reason. The difference between 5 inch drives and 8 inch drives is typically only a few hundred dollars at best. Here are some of the items to consider in deciding which size disk drive to specify:

1) PHYSICAL LIMITATIONS - If size of system must be held to an absolute minimum. (Lean toward 5 inch)

2) ACCESS TIME - 8 inch drives are SIGNIFICANTLY faster than 5 inch. (Lean toward 8 inch)

3) STORAGE CAPACITY - 8 inch DSDD drives store about 1.2 Mb, each 5 inch "quad density" stores about 630Kb each. (Lean toward 8 inch)

4) COST FACTORS - 5 inch drives cost less to manufacturer. (Lean toward 5 inch)

5) TRANSPORTABILITY - 5 inch has NO transport standard. (Lean HEAVILY toward 8 inch)

6) SOFTWARE VARIETY/AVAILABILITY: New software and updates are released in 8 inch format. Other formats are produced as the market for them dictates - sometimes never. (Lean toward 8 inch)

7) MAXIMUM BID RANGE: If other procurement restrictions are placed on the system, allowing 5 will open the bidding to more vendors. Disallowing 5 inch drives will eliminate perhaps 35 to 45% of the market today, but that still leaves hundreds. (Lean toward 5 inch)

Well, now that you're sea sick from leaning back and forth, I hope you are left leaning toward 8 inch, because:

RECOMMENDATION: Dual (two drives) DSDD, 8", soft-sectored, CP/M format (approx 2Mb total).

SYSTEM BUSS

Now we come to one of the only truely bonafide, formal micro industry standards - and yet the one requirement recommendation that can be manipulated the most.

A computer system can consist of a number of functional units, including the CPU, RAM, disk, I/O and other elements. There are basically two ways for a computer manufacturer to design a new computer. He can hardwire everything together, usually on a single board, or he can use the "buss" approach.

An electrical (signal) buss is any common medium on which signals can be

placed by one buss user and picked up by any others on the same buss. A TV transmission cable can be considered a simple buss, because a signal, originating at the antenna, can be tapped by anyone along the way who attaches to the cable. In the computer world, things are a little more complex. A buss can be a collection of many signal lines, over a hundred in some cases. There are three major catagories of busses presently in use in the microcomputer world - the system buss, the I/O buss, and the Local Area Network (LAN).

The LAN is normally used to connect complete computer systems together. All computers in the LAN can talk to any other computer, or to intelligent devices such as printers, file servers or external (phone) communications controllers. Most of the LAN's in use today operate over a single conductor coaxial cable - normally ordinary closed circuit TV (CCTV) cable. LAN's will be discussed at another time, however, since these systems are for the most part out of the performance and size range we are talking about right now. They also need the same level of planning and consideration that a large multi-user system needs. (Yes, there are CP/M based LAN's)

An I/O buss is used for communications between a computer and some kinds of external equipment (peripherals). The most popular I/O buss in use today is the IEEE-488 buss, also known as the General Purpose Interface Buss (GPIB) or the Hewlett-Packard Interface Buss (HPIB). This buss has been used in labratories and industry for many years to control scientific instruments, medical equipment and computer peripherals. All peripherals attached to the buss are 'Listeners'. Some devices are also 'Talkers'. A talker tells a listener to do something, and may pass some information to it. That listener may then become a talker, if necessary, and say something to another device. In microcomputers, the primary use of the IEEE-488 buss is to attach items such as printers, plotters, and sometimes even disk and tape drives. An I/O buss will usually consist of between 15 and 30 signal lines, and it has a limited physical range it can operate over (normally 50 to 100 feet). It costs significantly more to implement than some of the simpler interfaces such as the RS-232C (serial). An I/O buss can be added to many computers, but it should only be considered in the specification only if there will be a specific use for it. An RS-232C interface is used far more frequently for attaching items such as printers.

The final catagory of buss, the system buss, is used to tie together the various functional units inside a computer. The number of signal lines in a system buss will typically range from 50 to 150 lines. This collection of signal lines, along with the sockets which allow boards to be attached to them, are called the 'motherboard'. Each user of the buss, normally a functional unit such as a processor or memory board, resides on a board called a 'daughter' board. The principle is the same, however. One user of the buss talks, everyone else listens. If one user asks another user to do something, it does it. If the second board needs to talk to someone else, everyone else listens. A true system buss is distinguished from 'expansion' slots by the ability for any daughter board to be placed in any position of the motherboard and operate. The placement of a board is not dependent on hardware or software restrictions. A system buss can offer enormous benefits to the microcomputer user, but it has some potential drawbacks as well. Knowing when a buss is needed, and which buss to specify, can be quite important to your system's future.

There are many system busses in use today, and most of them are proprietary. A proprietary buss means that the company who originated the buss maintains absolute control over the buss specifications, and controls licensing for the use of it by others. A popular example of this type of buss is the LSI-11 buss used by Digital Equipment Corp (DEC). Some busses are controlled by the original company, but anyone can use it. The popular Multi-Buss originated by Intel Corp is an example of this type, and it is now under consideration for adoption as the IEEE-796 standard. The third type of system buss experienced such prolific growth that there was never any chance for any one manufacturer to lay claim to control. I refer of course to what is now the IEEE-696 buss, more commonly referred to as the "S-100".

The S-100 started as a chance selection on the part of an engineer at the MITS company as he was designing the Altair 8800 computer. For those of you who are new on the block, the Altair was the first micro computer to really make a splash in the personal, hobbiest and buisness market - in spite of what the Apple Corporation's advertising would have you believe. The engineer at MITS counted the number of pins he needed for his buss design, added a healthy fudge factor for future unused lines, and came up with a 100 pin selection. He had no idea what he would start that day. I learned fast about the advantages of the S-100. As an owner of one of the first Altairs, I soon found myself inundated with expansion possiblities. Everyone was competing to make their board the best possible for the least possible price. With every magazine issue, I watched the price of products fall as their capabilities rose. Board manufacturers could concentrate on the types of products they knew best. System houses and other integrators (read hobbiests) could choose between the very best and most cost effective boards to make up their system.

Unfortunately, some board manufacturers used certain undesignated lines of the buss for one special feature, and another used the same line for something else. Fortunately, very early in the game, all of these manufacturers got together and decided to insure their own survival by creating a formal buss standard around the S-100. Today's IEEE-696 buss not only remains compatible with nearly all of the older S-100 boards, but provides for incredible expansion and upgrade possiblties. For instance, (Novices may wish to skip or read lightly over the technical mumbo jumbo which follows) up to 16 parallel microprocessors can act as concurrent, prioritized buss masters. They will be able to operate at speeds of at least 8 MHz (possibly up to 12MHz) and access up to 16 million bytes of continious (NOT bank selected) RAM. In addition, there are 65,000 I/O ports available to communicate with the outside world. The possibilities are staggering. The S-100 protocols will allow mixtures of older 8 bit devices, such as memory boards, with newer 16 bit ones. Automatically. That is why you can literally remove your 8 bit processor and slap in a 16 bit one (or dual processor) and not have to replace your old memory and peripheral boards. It works! That comes under the catagory of investment protection. On top of all this, the S-100 is physically smaller, less complex, and cheaper to make than its only real micro competitor, the Multi-Buss.

Need some more incentive? There are an estimated half million S-100 systems in use today, with approximately half in dedicated industrial and process control applications and most of the rest in small buisness. Over 400 boards are available to choose from, made by hundreds of manufacturers. There are dozens of 8, 16 and 32 bit CPU's available on the S-100, including the 8080, 8085, 8088, 8086, 8035, 8039, 8040, 8048/49/50, Z-8, Z-80, Z-800, Z-

8080, 6800, 6802, 6809, 68000, 6502, TMS-9900, AMD-2901 Bit Slice, 2650, WD-16, Alpha Micro, Pascal Microengine, F-8 and the iAPX-432. These CPUs can come configured as stand alone processors, master/slave processors, distributed processors, concurrent processors, associative processors or auxiliary processors. The auxiliary processors can come configured as graphics processors, speech/music processors, floating point processors, test instrument processors, I/O processors or food processors (Just kidding!). There are some misguided (or greedy) souls who would maintain to their dying breath that the S-100 is limited in capability and can only be used effectively by the 8 bit 8080 processor for which it was originally designed. They are quite obviously wrong.

So, we have this really neat, super powerful capability. "How does it affect me?", you ask. As a potential purchaser of a computer system, you will be glad to know that you are only interested in some very simple, basic questions in your decision to include the S-100 in your specifications. The first and most important consideration is that of expansion. If you are planning to expand or upgrade your system beyond the normal addition of memory, disk and printers which most of today's systems support, then you need the S-100. If you are not sure about your future plans, you need it. If board level maintenance is critical, such as in vital shipboard applications, you should specify it. Weigh these factors against the slightly higher cost (usually $250 to $750) and some reduction in bid scope.

RECOMMEND: Designed to IEEE-696 (S-100) buss specifications.

OTHER FACTORS TO CONSIDER

There are several other items which should be considered prior to writing your final specification. These items, while not absoloutly essential to the operation of your system, should be included along with the other requirements we have developed. If an item will not be required immediately but may be required at some time in the future, then you will still want to include it as an 'expansion' capability. To leave out potential requirements may mean an inability to add them to your system at a later date. Some generalized guidance is provided here on some of these items.

PRINTERS

Most people will require a printer on their microcomputer system, and nearly all systems provide a means to attach a printer. I would recommend that allowances ALWAYS be made to attach a printer to your machine, even if you think you're certain you'll never need one. Within one month after you have your machine, a mandatory requirement for a printer will mysteriously appear. The ease with which you can add that printer could be an important factor to consider. Serial (RS-232C) interfaces are the most frequently used, and it is recommended that at least one port (a computer's "gateway" to the outside world) be free on you system for this purpose.

The printer itself should be capable of providing the speed you will need. You don't want an 80 Character Per Second (CPS) printer if you have to print out monthly inventories for thousands of items. On the other hand, a 200 CPS printer may be wasted if you only need monthly listings of a few hundred items, and a few short reports now and then.

You will almost always need a 132 column report format. The power of printed listings is to provide dense, meaningful information summaries. You may or may not need to get a wide carriage on your printer. In the Micro Team, we do virtually everything on regular sized (8 1/2" wide) paper, using compressed print to get our 132 columns. If you need a wide carriage, be sure to specify it. Whatever your volume and carriage size needs, you will also want to specify "bi-directional" and "logic seeking", which are two fancy terms that mean the printer figures out the most efficient way to print a new line based on where it stopped printing the previous one. It makes a BIG difference in the printer's speed, and it's available on virtually all modern printers.

To get an idea of what you may need, visit a local computer store and look at their low capacity printer (80 column, approximately 80 CPS), and their medium or high capacity printer (132 column, approximately 150 CPS). Try to guess how much printing you'll be doing, and see which speed printer will be most suitable. Most printers currently range from about $750 for the low capacity to about $1800 for the high capacity. A high capacity printer can usually be justified for most serious applications.

Most of the same advice applies to letter quality printers, except that speeds typically range from 25 to 55 CPS and price ranges from $1800 to $3000. Obviously, letter quality printers are not suitable for report generation in even medium scale situations. Most systems will NOT need a letter quality printer unless high quality outside correspondence is required, or OCR forms must be generated. Modern high quality dot matrix printers can provide excellent print quality, more than adequate for nearly all Navy correspondence needs. Again, look at the possible machines at a local computer store. You may even want to specify a fast dot matrix printer for reports, internal correspondence and drafts, and a very low speed (25 CPS), low cost letter quality printer for outside correspondence and/or OCR generation.

COMMUNICATIONS

A communications capability is available on nearly all CP/M based machines. You may wish to use this capability, so that documents, databases or other information may be easily transferred to another location. There are many times when you will wish to exchange information, programs or data with your counterparts in other areas of the country (or world), or even people you don't even know right now. This capability also allows you to use the information and conferencing facilities available on the dial-up Navy Bulletin Board System (NBBS).

There are many different communications protocols available under CP/M. (A protocol is the way in which one machine talks to another) The 'standard' CP/M protocol uses a public domain (free) program called 'MODEM', and normally uses a 300 baud dial-up modem. More specifically, (again, new folk may skip this) the modem talks using a 300 baud Bell 103 compatible, asynchronous

format with 8 data bits, one or two stop bits, and no parity. The standard protocol sends "packages" of 128 characters each with a two character "checksum", and ACK/NAK verification. The public domain MODEM program is available from the NARDAC Micro Team at no charge, although you may have to modify it for your specific installation.

Some very excellent commercial programs are also available, although most of them will use proprietary protocols. These commercial packages allow communications with "mainframe" and popular minicomputers, as well as other microcomputers. They will also allow communications with other computers using popular "synchronous" protocols (3780 ("bi-sync"),X.25, etc). Both ends of a communications link would have to buy a copy of any commercial (proprietary) package. Compatability can still be maintained with every one else through the public domain packages, while taking advantage of the somewhat more friendly commercial packages. This is called having your cake and eating it too.

As a minimum, you should get an inexpensive 300 baud acoustic modem which you can attach to one of your serial I/O ports. For added convenience and reliability, you could get one of the newer "direct connect" modems which plug directly into a standard modular phone jack. This eliminates all outside noise and results in less "connection hassle". For the communications connoisseur, auto dial/auto answer units are available which will allow your computer to call another computer automatically under program control. The other computer can answer automatically, and the two can conduct business from afar without any human intervention. Anyone ever see "The Colossal Project"? Who needs dumb humans anyway?

SPECIAL POWER CONSIDERATIONS

Many operating environments are not as electrically ideal as we would like them to be. As a matter fact, most would give computer equipment designers heart failure. If the power where the computer is to be used is trashy (noise, spikes, surges, dips, etc), but does not fail completely very often, then you should specify some sort of line conditioning. This problem is especially prevalent aboard ships. The easiest solution is to specify that the power supply of the computer be of the Constant Voltage Transformer (CVT) variety. This kind of feature is available for many S-100 buss systems, and results in only a slight increase in cost. If the system is non S-100, then you will have to specify a standalone power conditioner. A CVT power supply or conditioner will normally keep your system operational from 95 to 140 volts. Take heed that many of the inexpensive "line filters" will only take care of line noise and spikes. They do NOTHING for surges and dips. If you have problems with complete power loss, you may wish to obtain an Uninterruptable Power Supply (UPS). A UPS is especially advisable for shipboard applications where a hard disk is used. It will normally cost between $2500 and $4000, depending on its capacity, but it may pay for itself with the cost of a single saved hard disk drive.

ENVIRONMENTAL CONSIDERATIONS

Now we've all been at various Navy commands. We know how spotlessly clean they are, and how ideal their temperatures and humidity are. But for

those few cases where dust, temperature and humidity are a problem, you may wish to specify the conditions under which the system will be expected to operate. If the system will be in a dirty or dusty environment, an air filter should be provided as opposed to the more frequent method of having an open grill. Most micros today will operate reliably within a range of about +50 to +90 degrees Farenhenheit and with humidity from 20 to 80 percent (non-condensing). Environmental ranges outside of this are usually not considered acceptable working conditions. If extended operation of the system outside these ranges is expected, be sure to include the exact requirements in your specification.

Vibration is the last condition to worry about. Continious vibration, such as that which exists in most shipboard spaces during normal steaming, is no problem. Sudden shocks, however, can be dangerous to the equipment. If the system is to be installed in locations near cannons, catapults and other shock producing machinery, the equipment should be mounted on standard Navy shock mounts. A shipyard can probably calculate the proper size and arrangement of the support mechanism.

MAINTENANCE

Microcomputer systems are composed of electrical and mechanical devices. Both types of devices WILL eventually fail. Within the first 30 to 90 days, you can expect problems with most systems. This time is referred to as the "burn in". Once this burn in time is over with, the reliability of the system will be extremely good, averaging about 2.2 years between failures. This is assuming proper environmental conditions have been maintained and that preventive maintenance (dusting, contact cleaning, etc) have been performed. This also does not include such minor problems as dropping sodas down keyboards and pulling boards out for cleaning with the power on, and similar actions (Believe me, it happens). When a failure does occur after burn in, it will nearly always be due to heat, dirty contacts, or damage from electrical power fluctuations. Mechanical devices (disk drives, printers, etc) are far more likely to fail than purely electrical ones, but these have also proven to be quite reliable.

You should always insure that you will receive full maintenance, under warranty, for the burn in period. Most vendors will give a 30 to 90 day guarantee. If they don't, you may want to consider an "extended" warranty out to perhaps 6 months. After that, it seldom pays to obtain a maintenance contract.

Service contracts average between 1.5 and 2 percent of the total system cost per month. Many microcomputers today are serviced through "third party" service organization, such as TRW, Sorbus, General Electric, and others. If you choose a service contract, be sure to specify or be aware of the response time. Although some contracts are available which guarantee repairs in as little as 4 hours, most specify repairs will be made within 48 hours. Some specify 7 working days, and some don't specify anything at all! And remember, even the best service contract won't do you one bit of good if you're floating around on the ocean!

Another solution, especially for shipboard applications, is to stock a spare board/critical part set. This allows immediate repairs by

onboard technicians. Such "spares kits" can be easily put together by the vendor. Training should also be provided for, in this case, either by the local vendor for from the factory. Only board level maintenance is recommended, unless a large, talented pool of microelectronic technical expertise is available (Such as tenders and carriers would have).

If the application is not so critical, an "as required" method can be used. This means that you just keep running your system until something finally fails, then call your vendor in to fix it. Bear in mind that an "as required" customer will fall at the bottom of the priority list for a vendors service department - his contract customers must always come first. Sometimes it's even feasible to work directly with the factory. Many companies can trouble shoot your computer over the phone, and will have the repair part shipped by express overnight delivery. If an exchange plan is desired, the bad component is returned with a flat fee - usually $50 to $150, depending on the item. Either of these two alternatives should be provided for with a BPA (Blanket Purchase Agreement) or some other rapid procurement vehicle, so that repairs are not hampered by unexpected paperwork. Most repairs will fit easily under a $500 threshold.

TRAINING

In addition to the the user and repair manuals that come with your system, you may wish to specify that training be included with the system. This may be in the form of training from the dealer or from the many available packages which utilize Computer Assisted Instruction (CAI). The more popular word processing, data management and spreadsheet software have these packages available, as well as for CP/M itself. They are considered a good way to start off beginning users. Dealers or factories would have to administer any hardware or maintainence training should that be required. On the Job Training (OJT) is also a very effective method of training in the micro world, due to the fast operator learning curve of today's software.

SOME FINAL THOUGHTS

Remember that in today's market, much of the well advertised equipment is targeted for home consumers and professional managers. Have you ever wondered why "personal" computer systems cost so little compared to "business" systems?

PCs are designed to have an exceptionally low entry cost. The typical consumer could simply not afford to plop down several thousand dollars in one big lump. The typical consumer is also far more concerned with immediate (short term) cost than with performance. To satisfy those users who outgrow their initial machine, many PCs will allow considerable expansion, allowing larger drives, more memory, different displays, sometimes even different processors. There is normally no profit or efficiency motive to justify a performance increase. Those PCs which are used for profit are normally used in small, desktop type applications for which the PC is most effective.

A business, or Navy command for that matter, IS driven by efficiency. If a computer is justifiable at all, it will usually (easily) justify those thousands of dollars. A complete system will consist of the basic machine (CPU, full memory, I/O ports, etc), two floppy disk drives, 132 column

printer, modem, and a CP/M compatible operating system. Be prepared to commit an absolute MINIMUM of $5000 for your complete system, and expect about $8000 or more for a more full-featured system with a letter quality printer. If anyone tries to sell you this configuration for more than $10K, start walking. Keep in mind that if your system must be leased (i.e. OM & N funds), you may have to select from more costly systems, since very few low overhead companies can afford to lease to the government. In this case, the features MAY exceed the $10K threshold noted. Justify your funds (estimate $8000) AND the source of the funds BEFORE you proceed.

Five to eight thousand dollars? That's outrageous! That PC only cost about a thousand! Why not just buy that PC and just expand it? Do this little experiment. (NOTE: Be sure to use identical specs (or brands) if possible, and use list prices)

1) Start with the cost of a "basic" personal computer.

2) Add the cost of features required to bring that system up to the recommendations in this article. (Note any time you are prevented from doing so because of a lack of the PC's flexability)

   * CP/M 2.2 DOS
   * Z-80 CPU (4 MHz)
   * 56K TPA
   * 80 x 24 CRT display (12", green)
   * 'QWERTY' keyboard w/shift and control key (May have fun with this on some PCs!)
   * Dual DSDD 8" soft sectored disk drives (approximately 2Mb total)
   * 132 Column Printer
   * Modem

3) Assuming that you have successfully expanded the PC to include these features, compare the cost of the "expanded" PC with the cost of a buisness system with identical features and consider the following:

   * How much savings is there as a function of the total system cost? There will typically be less than a 15% difference, IF ANY!
   * Its power supply is being taxed to its limits, and overheating will very likely become a major concern. An outboard cooling fan will probably be needed.
   * The expanded system is probably filled to near capacity. You have little, if no, room to grow for other expansions. (A typical business system is just beginning)

Sales people, well meaning though they they may be, will often attempt to sell things which are not really in your best interest. Let me re-emphasize that I am NOT "anti personal computer". We use PCs at the Micro Evaluation Team, and we often recommend them to fulfill numerous applications. But different machines exist to serve different needs. You don't use a row boat to go water skiing, even if the rowboat salesman tells you its no problem to add a large engine, steering mechanism and a reenforced transom. On the other hand, a ski boat is a costly investment if all you want to do is fish.

All salesmen are NOT out to get you. But you must know what to ask, and be familiar enough with the what you are trying to buy not to be misled (A dog will byte you if you it thinks you are scared). Remember that for a salesman, benefits sell, features don't. A mediocre feature is often changed into a rip-roaring benefit. The mediocre feature of an "8080 processor" can become the benefit "Uses proven, reliable microprocessor technology". Get the idea? Change benefits back into features (also known as "plain English") before comparing systems.

Most vendors will be very helpful - they want your business and your references. And unless you have access to a trained technical staff, you will need the services of a good computer store or system house to assemble and test your system, install it, train your workers and provide service. Do NOT buy through mail order until you have considered ALL costs. Mail order firms market at minimum mark up, but provide NO after sales assistance. For many Navy commands, that local friend can be a vital link in getting your operation off the ground.

A microcomputer can be one of the biggest assets you will ever have. They are hard at work right now in all aspects of everyday life. Don't let the buzz words you hear get to you. I recommend you buy or subscribe to some of the "micro magazines", such as Interface Age, InfoWorld, Kilobaud, BYTE (for the technical types), and Personal Computing (for the novice types). Their reviews and articles will be of immense help to you. Naturally, the NARDAC Norfolk Micro Evaluation Team is always ready to assist you with any problems you may have, and our Navy Bulletin Board System is ready to answer technical questions. Hope I've been able to help you. How about letting me know what kind of a system you get, and how it works out for you? We always have plenty of coffee on the fire.